

InvestigationOrganizer: The Development and Testing of a Web-based Tool to Support Mishap Investigations

Robert E. Carvalho
NASA Ames Research Center
MS 269-1
Moffett Field, CA 94035
650-604-3593
Robert.E.Carvalho@nasa.gov

James Williams
NASA Ames Research Center
MS 213-4
Moffett Field, CA 94035
650-604-6377
James.F.Williams@nasa.gov

Ian Sturken
QSS Group, Inc.
MS 269-1
Moffett Field, CA 94035
650-604-0722
isturken@email.arc.nasa.gov

Richard Keller
NASA Ames Research Center
MS 269-2
Moffett Field, CA 94035
650-604-3388
Richard.M.Keller@nasa.gov

Tina Panontin
NASA Ames Research Center
MS 213-10
Moffett Field, CA 94035
650-604-6757
Tina.L.Panontin@nasa.gov

ABSTRACT¹

InvestigationOrganizer (IO) is a collaborative web-based system designed to support the conduct of mishap investigations. IO provides a common repository for a wide range of mishap related information, and allows investigators to make explicit, shared, and meaningful links between evidence, causal models, findings and recommendations. It integrates the functionality of a database, a common document repository, a semantic knowledge network, a rule-based inference engine, and causal modeling and visualization. Thus far, IO has been used to support four mishap investigations within NASA, ranging from a small property damage case to the loss of the Space Shuttle Columbia. This paper describes how the functionality of IO supports mishap investigations and the lessons learned from the experience of supporting two of the NASA mishap investigations: the Columbia Accident Investigation and the CONTOUR Loss Investigation.

1. THE NEED FOR INVESTIGATION SUPPORT TOOLS

Currently, mishap investigation teams face several challenges when determining root causes and improving safety: the increasing complexity of systems, barriers to efficient information distribution, and limited methods for preserving and standardizing investigation results.

The increasing complexity of systems is demanding more time and effort from investigation teams working to adequately identify failure causes. This can create a need for more complex investigation models, including hybrid models, which can track in detail the progress investigators are making on individual model elements or on hypotheses. During an investigation, teams monitor

requests for information, investigation actions, and the supporting or refuting evidence for each part of a model. Additionally, a wide range of information types -- drawings, documents, interview recordings, video images, and telemetry, among others -- is being brought together to provide the full picture of what happened.

Because investigation team members and the investigation sites are often geographically distributed, sharing data quickly and securely among investigators is another challenge. Although e-mail is frequently used, sharing all documents via email is often impossible or prohibited because some email servers prevent sending large documents or because unencrypted email is insecure. Investigators sometimes resort daily to overnight shipping of boxes of CDs, photographs, and paper documents. Meanwhile, because causal models are typically developed on white boards or paper charts on walls, teams often cannot share them off-site.

Disparate approaches to investigations dilute the value and distribution of lessons learned. In some organizations, the data and notes generated during an investigation are stored in only the lead investigator's notebook or computer. There is little preservation of relationships between pieces of information, even in standard databases. In subsequent investigations, this can make gathering information from similar cases impossible. If a safety organization attempts to review older mishaps for missed precursors or other hazard trends, varying levels of detail in surviving analysis notes can make conducting detailed examination difficult.

2. DESCRIPTION OF INVESTIGATIONORGANIZER FUNCTIONALITY

InvestigationOrganizer (IO) is a web-based tool that was designed to help investigation teams overcome some of these challenges. IO provides a common web-accessible repository for a wide range of mishap related information,

¹

and allows investigators to make and share explicit, meaningful links within causal models, which show how one event may have caused another event, and between information such as evidence, findings and recommendations. IO integrates the functionality of a web database, a shared document repository, a semantic knowledge network, a rule-based inference engine, and causal models. The semantic knowledge network [1] allows users to create different types of items, which contain information associated with a thing, a place or an

idea, such as a person, a piece of evidence, a system, an action item, or a causal model. Most importantly, the semantic network allows users to create meaningful links between these items; for example, an investigator might collect a piece of evidence or a document that supports a hypothesis (see Figure 1). These meaningful links allow users to explicitly share their reasoning as they conduct an investigation, "connecting the dots" -- from evidence through analysis and modeling, and into findings and recommendations.

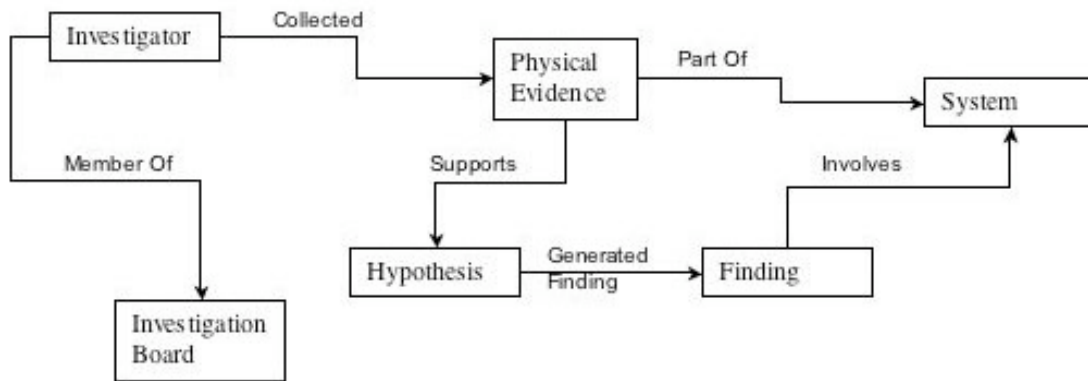


Figure 1. Semantic Knowledge Network

In the early phases of an investigation, investigators often gather information from many locations. IO serves as a common repository for digital data, enabling instant and secure information distribution to all team members as an investigator makes it available. Speedy distribution to team members reduces the chance that more than one investigator will waste time chasing down the same pieces of information.

As the investigation moves into the analysis phase, IO allows investigation teams to develop, view, and share causal models (such as fault trees or event sequences) and hypotheses that teams will use to ensure they are considering all possible causes of the mishap. Most importantly, the semantic network allows users to create and share meaningful links between these models, the analyses, and the evidence (see Figure 2).

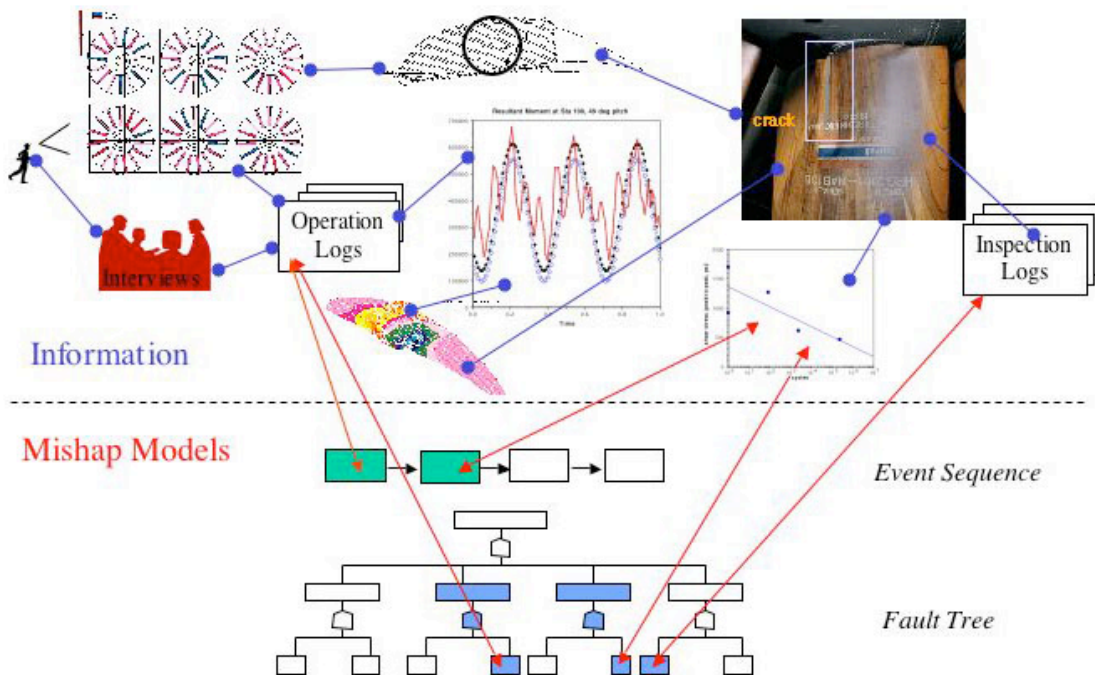


Figure 2. Linking Information Items to Causal Models

The establishment of meaningful relationships between the data and analysis information provides organizational memory for investigators to preserve their reasoning in the conclusions. For example, say a photo supports an element of a fault tree, the element of the fault tree generates an action item, and the action item produces a finding --this memory can be valuable to teams implementing safety recommendations for the investigators, and to future investigation and safety teams and reviews.

3. HISTORY OF THE DEVELOPMENT AND DEPLOYMENT

InvestigationOrganizer (IO) was created and funded by the Engineering for Complex Systems Program within NASA. In February 2002, IO development began, leveraging off an existing semantic network system prototype that had been developed to support scientific investigation teams [2]. The initial IO prototype included a new ontology, or structured language, for mishap investigations and new prototype viewers for different models. The ontology describes the types of information items that users can create in the system (e.g. people, vehicles/systems, evidence, causal models, hypotheses, findings, and recommendations) and the available relations between them (e.g. person authors document, evidence supports hypothesis, document is response to request for information). The prototype viewers were

specialized for viewing fault trees and event sequences. By June, the prototype was complete, and was tested with data from a small investigation that had been conducted several years earlier. IO was used:

- On July 27, 2002, during an air show at Moffett Field, CA, a minor mishap damaged windows on the base and caused superficial injuries to two air show staff members. IO was used in the investigation as a beta test of the system. The use of IO on this investigation involved 5 users and more than 200 information items.
- On August 15, 2002, the Comet Nucleus Tour (CONTOUR) Spacecraft was lost while firing its main engine to move out of Earth's orbit on its way to intercept a comet. The IO team was invited to support the investigation, again as a beta test. The use of IO on this investigation involved 25 users and more than 800 information items, including a 125-node fault tree.
- When the Space Shuttle Columbia was lost on February 1, 2003, IO was requested for a portion of the Columbia Accident Investigation Board (CAIB) in their investigation into the direct or proximate cause of the accident. The use of IO on this investigation involved more than 100 users and 4800 information items, including a 400-node fault tree and a 200-node event sequence.

- On June 26, 2003, the Helios solar-powered Unmanned Aerial Vehicle broke up in flight near Hawaii, and IO was asked to help support that investigation. The use of IO on this investigation involved 23 users and more than 1400 information items, including a 168-node fault tree.

Based on the increase in demand for the tool, in June, 2003 NASA began transferring the technology to Xerox under a NASA Space Act Agreement to develop a commercial-grade system for use not only by NASA, but also by other government agencies and industry.

4. CONTOUR MISHAP INVESTIGATION

4.1 Overview

The unmanned CONTOUR spacecraft was launched from Cape Canaveral on July 3, 2002. Its mission was to investigate two comets by taking pictures and performing various analyses. Contact with the spacecraft was lost as it attempted to leave Earth's orbit on August 15, 2002. Pieces of debris were subsequently spotted along the spacecraft's trajectory, suggesting the loss of the vehicle.

During the week of August 21, an investigation board convened at NASA Headquarters in Washington, D.C., including 5 board members and investigation support staff. The board and its staff would work from their offices throughout the US, at headquarters and in states as distant as Alabama and California; the need for a tool that would enable them to organize and share information from a distance was apparent. Since a member had seen a demonstration of IO a few months before, the board asked to use IO.

By August 26, IO representatives traveled to NASA HQ. Since development of IO was not complete, they laid out a plan for the investigation board to use the tool in "shadow" mode. In this mode, board members would use the system to access data, but the IO support staff would handle data entry. (Some board members eventually developed enough experience to enter their own information.) One IO team member was designated an observer, to gather information during board functions and meetings then enter the information into IO. Board members would be set up with accounts and trained to use IO by the end of August 2002.

In the first month of the investigation, a large amount of information had to be created in IO. A great deal of this was investigation infrastructure-- information that would provide the foundations for organizing investigation information in the days ahead, such as a model of the CONTOUR spacecraft, design information, personnel

descriptions, and review documents. To handle the high volume of information to be created in IO, several team members were deployed to NASA HQ backed by IO team members at NASA Ames Research Center. As the investigation progressed and the stream of information lessened, the number of personnel needed in D.C. and at Ames decreased. By January 2003, all IO work was being done at Ames Research Center, with the Investigation Board communicating with the IO team via email and telephone.

The CONTOUR final report was published May 31, 2003 [3]. The final task for the IO team was creating an archive CDROM of all CONTOUR information.

4.2 IO Usage During Contour

The Contour mishap board used IO in several capacities including information sharing/organization, causal analysis, managing workflow, and archiving.

IO was used primarily for organizing and sharing information. Throughout the investigation, the team added and linked information in IO. Documents, photos and schematics were all loaded into IO and were available anywhere to anyone with an authorized IO account and a web browser. The ubiquitous nature of IO, the board said, was one of its greatest assets. Another capability enabled the team to add placeholders for hard copy information that wasn't available in digital form. These documents were generally located at NASA headquarters in a secure room. The placeholders provided a mechanism by which investigators could determine if a document was available at HQ, even if it was not in IO.

To keep the investigation team up to date on the information going into IO, daily status updates provided a description of the information and a web pointer to the information in IO. Board members could quickly scan these reports, click on the links and retrieve the newly added information.

E-mail for the investigation was archived in IO and linked to appropriate nodes. For example, if an email discussed a particular design document, the email would be linked to the design and an easy reference would be provided in the subject of the email. The design document would also be linked to the system it described, providing a cross-reference (Figure 3). This way, if an investigator remembered that a particular document came in an e-mail from a person on a particular day, or if they only knew that it related to a particular system, they could search for it based on the information they had available.




Figure 3. Email Links in IO

Another heavily used feature in IO during the CONTOUR investigation was the fault tree (Figure 4). Investigators developed a fault tree in IO with the support of the IO team, assigning nodes to investigators and tracking progress in the investigation. As the investigators turned up information to support or refute the part of the fault tree they were working on, they could upload the information, and link it to the node in question. Each node had a color code indicating the status of the node. For example, grey nodes were "Closed – Not Credible," orange nodes were "Needs Analysis." Members of the board could therefore determine the state of the investigation at a glance by viewing the fault tree. In this way, the course of the investigation could be watched by

viewing the tree; in the early stages of the investigation most of the nodes were "Needs Analysis." As the investigation proceeded, nodes slowly moved to "Closed – Not Credible." At the end of the investigation, one branch of the tree was left: "Closed – Credible." If one were to look at the tree today, one could quickly surmise the cause of the mishap and find the key pieces of information that supported this conclusion, by following the Closed – Credible branch in the tree. The state of the investigation was not the only information one could get from the fault tree. Individual investigators would know which nodes of the tree they were responsible for by viewing the nodes that were linked directly to them.

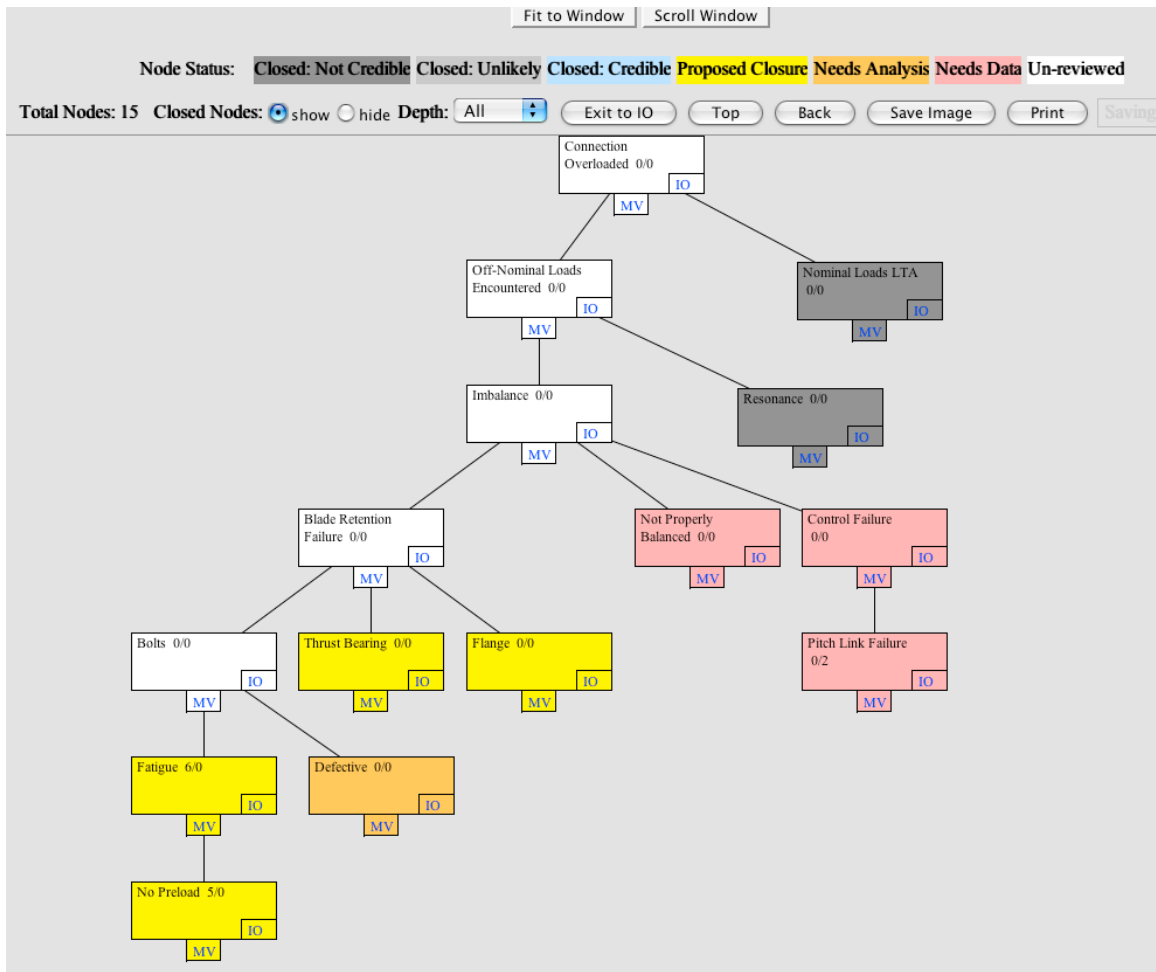


Figure 4. Example of Fault Tree in IO

As can be seen in Figure 4, the fault tree viewer in IO provides information and navigation to help the investigator. The colors clearly show the status of each node. The numbers in the boxes show how many pieces of evidence support or refute a particular fault tree node. (e.g. Pitch Link Failure has two pieces of evidence that refute it.) To find out more information about a particular fault tree node, the user can click on the IO button to return to the standard IO interface on that node and see who it is assigned to and what evidence is supporting or

refuting.

The Contour board relied on IO for workflow management. The investigation board would assign action items or requests for information (RFIs) to investigators. Investigators would perform the requested action or obtain the information and provide the results to the board. Once the action or RFI was satisfied it would be closed. This process (shown in Figure 5) was modeled in IO.

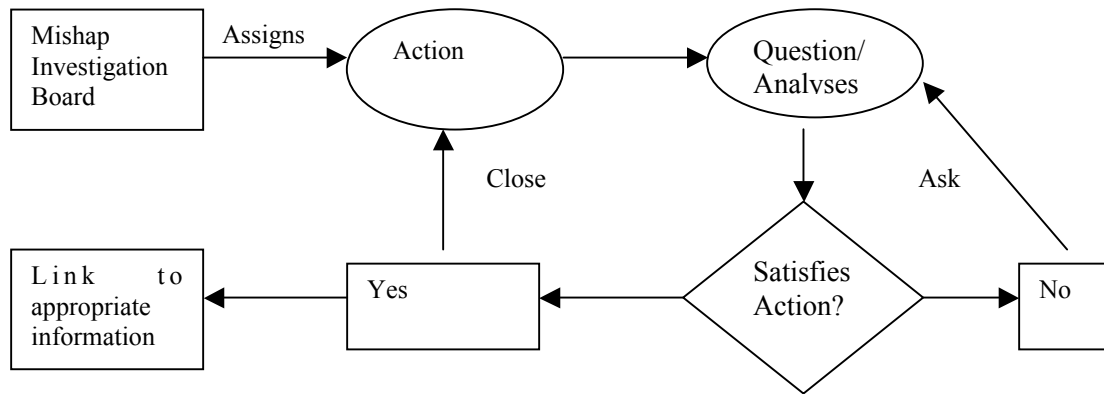


Figure 5. Workflow for Action Items in IO

IO provided a means by which the actions could be tracked. The added bonus was that the information resulting from an action could be found linked to the action within IO. This didn't always work, though; sometimes actions would be satisfied via email conversations that were difficult to incorporate fully into IO if users did not copy the archive on all of the e-mails.

In many cases, investigations are archived in some informal way, perhaps in a box with some documents and CDROMS. IO provided an easy mechanism by which a "snapshot" archive was made of the CONTOUR investigation at its conclusion and was placed on CDROM. These CDs preserve not only the information but also the reasoning within the explicit links between that information. This is done through a static html preservation of all the pages of IO for the Contour.

5. COLUMBIA ACCIDENT INVESTIGATION

5.1 Overview

On February 1, 2003, the Columbia Space Shuttle disintegrated upon re-entry into the Earth's atmosphere, resulting in the loss of 7 astronauts. The majority of the Columbia fallout debris ended up scattered throughout Texas and Louisiana.

The IO team was dispatched on February 3 to Barksdale Air Force Base in Louisiana to determine if IO could be utilized to help catalog the debris. Within a few days, the team concluded that it would not be able to meet the debris categorization needs of the NASA Accident Investigation Team (NAIT), mostly because, at the time, the information could not be uploaded fast enough, and because reports for the investigation could not be generated. These limitations existed because the system was still a prototype under development. Additionally, there had been so much debris collected already that the retroactive adoption of a new system would have been too burdensome.

Approximately 10 days after the loss of the shuttle, the IO team was asked to present the IO tool for potential use in supporting the Columbia Accident Investigation Board (CAIB) in Houston. During this time in Houston, two other software programs were presented to the Board as potentials for helping to organize the information and data that would be generated by the Board.

Initially, IO was chosen to be used as an investigation support software by one group of 4 board members and their support staff. This group was the "Direct Cause" group, trying to determine the technical reason for the loss of the shuttle. This was about 10 investigators. IO would be used primarily to link key pieces of evidence to various hypotheses. Because IO was a prototype system and not fully developed, the investigators needed to have three IO tool support staff present.

Another software called Process Based Mission Assurance (PBMA) was used to catalogue all of the NASA generated data for the accident, so the board began using it as the source for obtaining information from NASA. Requests for information were made to NASA and then processed using this software. This repository was accessible by board permitted people only.

About two months into the investigation, the Department of Justice (DOJ) dispatched a team of people to capture all information for archival purposes of the investigation. Since neither PBMA or IO were certified to capture and catalogue every piece of information collected and generated by the CAIB for permanent archival into the national archives, the CAIB had to employ the services of the DOJ to capture this. This team came in and added another level of information technology, creating an additional level of complexity. Whereas all three software teams understood the values of what they were doing, among the investigators was a perception that too many IT systems added requirements that slowed them down. This perception discouraged some people from actually providing data, and caused difficulty in the

addition of information in the various tools.

With about two weeks of accident data already collected, the IO team began to backfill available information into the tool. The investigation infrastructure developed, including a high-level system decomposition, but no investigation data had been added. Finally, the IO team began a single thread of the accident investigation with one key investigator helping to structure his information into the tool. This investigator was the single thread that was needed to begin demonstration of the intelligent value of IO. Within a week, this investigator's work was presented to the sub team of the board, and more people in this group began to provide data. Initially, this data came from the board members and investigators. Eventually, the NAIT investigators became comfortable enough to allow the IO team to start extracting information from the PBMA database.

Three weeks into the investigation IO was demonstrated to the entire CAIB. Based on this, the Board mandated that IO be used for all the groups in the CAIB investigation and that accounts be created for the entire CAIB team. The IO team went from supporting 10 people to supporting about 100 people and having to handle significantly more information. The expansion to supporting the whole board resulted in an immediate need for additional support staff. The lack of trained users among the investigators combined with the late request for use from the board resulted in a huge time lag and backfill of data that continued throughout the majority of the investigation. IO team members were assigned to each group within the investigation with several floating and one integrating the data across groups. As a result, the IO team members, in essence, became integral members of the groups and began providing investigation support by creating digital documents through scanning, helping to organize and categorize action items and related information, tracking document availability, and working with investigators to get information that they needed. This approach allowed the IO team to gain invaluable experience and allowed IO to capture the majority of the investigation process and information that the CAIB was performing. The IO team became a support staff for the investigation team and tried to facilitate the data organization and retrieval processes for the team.

Over the next month, a formal tracking system for internal investigation actions in the Columbia investigation was developed and implemented in the IO. This was based on the system developed for Contour, but reworked and expanded to include links to findings and recommendations. This action item-driven process proved to be a highly dynamic and valuable investigation technique in this accident, and one of the primary uses of IO. Individual investigators would be assigned action items, would gather background facts (evidence), and

would produce findings and recommendations to be submitted to the board. The board used IO to track the progress on the hundreds of action items in progress.

Also, within this timeframe, the CAIB developed an accident investigation "storyboard" as a unique causal model and method for ensuring that the investigation covered all possible avenues. This storyboard required the development of a different model in IO. The flexible nature of IO allowed this new model to be developed very quickly. The tracking of the investigation's progress against the storyboard was the other primary use of IO by the CAIB.

In early May, the CAIB moved to Washington, DC. By this point, the CAIB writers had already outlined the report and were no longer using IO to perform the integration of information. Most of this was because of inadequate report generation capabilities in IO. By the end of June 2003, the IO team finished its support efforts for the CAIB. IO had collected most of the data being used in the investigation report. The final report for the Columbia was released in August 2003 [4].

5.2 IO Usage During Columbia

During the course of the investigation, IO was never fully utilized. Although there was an understanding of the tool's capabilities, its late introduction, its user interface difficulties, and its lack of reporting capabilities contributed to this. However, aspects of the tool were used, and as a result the Columbia Accident Investigation approach and investigation process are now captured within the tool.

Once again, the primary use for IO was in the organization and sharing of information. The CAIB would make requests for information, and when these request were filled, the documents would be added into IO. On occasion, the IO team member assigned to the group would make an investigator aware of interesting information that someone else had added. Even though the CAIB was located in one building, the size of the investigation team and the amount of travel required to conduct the investigation made the use of shared data systems a necessity.

The sequence of events leading to the failure was preserved in the IO tool. This information was provided by the NAIT in Microsoft PowerPoint and other graphical formats. The IO team reproduced it in IO and linked it to systems and evidence, allowing investigators to see more information behind the sequence.

The fault tree was also used during the investigation. The fault tree was maintained and tracked by one board member and one assigned investigator. The fault tree followed was the same fault tree used by the NASA investigation team. This permitted the board to follow

and verify the fault tree closures independently of NASA.

The CAIB heavily used the action item tracking and linking features. Each action became a mini-investigation involving the specifics of the topic of the action. An action item would have background facts, findings, analysis, observations, and recommendations, as well as the appropriate references or documents. Throughout the lifetime of an action item, the appropriate information structure would be made. The key was the ability to use the action item feature to justify the findings and recommendations that were made during the course of the investigation. Action items would have background information and facts that resulted in observations and findings that led to recommendations. The action item would be supported with reference documents/evidence, and connected to the storyboard.

A unique challenge in supporting the CAIB was the development of a new cause model withing IO. The CAIB developed a modified cause model that was called “the storyboard.” The storyboard was a hybrid fault tree and hypothetical sequence. The flexible modeling capabilities of IO allowed even this unique hybrid model to be integrated very quickly into the system.

6. LESSONS LEARNED

Through the experience of using IO in live investigations, the IO team has learned how the tool can best be used, validated the approach and models for data collection, sharing, and analysis, and obtained requirements for future enhancements. For this paper, these lessons are divided into features that worked well, those that improved during the investigations, and recommendations for future use.

6.1 Worked Well

Not all of the investigations had distributed teams, and the size of the investigation teams varied significantly, from 5 to 100. However, even in the case in which all of the team members worked at the same site, IO greatly enhanced the sharing of data between team members, even within the same building. As the size of the teams became larger, having a good mechanism for sharing information about particular aspects of the investigation became even more critical. The capability in IO to link information to the causal model and to the subsystems or components made it much easier for different sub teams in the investigation team to see information being gathered that might relate to their own work.

The capability within IO to model the system decomposition and link important information directly to the system elements at any level made it easier for the investigators to see how all the pieces fit together and to find information about a particular subsystem.

The investigations took a long time. The shortest investigation time span was still seven months, and one was longer than a year. The ability of IO to preserve information and reasoning in the linking of the information items meant that investigators could more easily see what work was being done in different areas of the investigation. Other agencies or companies, however, may conduct much quicker investigations, and this could impact how IO is used.

6.2 Improved during investigations

All the investigations used an RFI and action item process with investigators, system developers, and operators. This was one of the primary ways of gathering information. Initially, IO included only limited capability in this regard. As IO has evolved, the capability to track and manage RFIs and action items has improved. Future work with IO (see section 7) will continue to improve and automate this aspect of investigations.

Full text search: The search function in IO was fairly primitive. It could only search the titles and user-entered key words of information items. This meant that the contents of documents could not be searched. It also could not perform the more sophisticated Boolean searches that are available in most modern search tools on the web. We addressed this by adding both the capability for Boolean searches as well as a more powerful search engine that could search the contents of the uploaded documents themselves.

Ontological Complexity: It turned out that the number of types of information and relations between them in IO were too great. This caused confusion among the investigators. The ontology has since been revisited and pared down to a smaller number of classes and relations.

The data about the systems involved in the mishaps came in a variety of forms. Some were electronic, some were paper, and some were mixed. However, within IO, only a subset of data regarding the critical subsystems and components were imported while the rest of the system was modeled at a higher level. This was done to reduce the amount of labor to input the full documentation. Research is continuing to create an interface between IO and existing design repositories to simplify this process. Additionally, access to better scanning tools would make converting paper documentation for access by the distributed team easier.

6.3 Recommendations

Two of the investigations had significant amounts of physical evidence, one had a little physical evidence, and the fourth had none. However, even though IO can support the tracking of physical evidence in investigations, this capability has not been exercised since IO was not brought in at the beginning of any

investigation. By the time IO had been brought into the investigations, other systems were already established for tracking physical evidence, or most of the debris collection had already been completed.

IO faced some problems and resistance in each investigation. In the high-pressured environment of a mishap investigation, investigators do not have time to learn a new software tool. IO's unique interface and data structures made it particularly difficult to learn. Also, as IO was brought in after the start of the investigations, a large amount of data was there to backfill, which required some investigator time. A more ideal application of IO would be for a mishap investigation team to be familiar with the tool before they begin an investigation. Additionally, a commercial quality interface would make the system easier to understand.

Report Generation: The investigators who used IO were most disappointed with its lack of ability to generate concise, meaningful reports about the information and links within the system. Given the interface to the system, it is difficult to get a big-picture view of the investigation outside of particular models like the fault tree. The future releases of IO are expected to address this critical gap.

7. FUTURE WORK

Currently, NASA is working with Xerox, Inc. under a Space Act Agreement to develop a commercial version of IO that is intended for use by government agencies and companies in a variety of industries. The initial version of this system will be available before the end of 2004. NASA is also continuing research into several major system enhancements to be included in future releases:

Navigability: Most of computer users are used to a hierarchical view of information using a folder scheme. Information in IO is organized as a network or graph rather than a hierarchy. This requires a paradigm shift for an IO user. In addition, as users follow the paths to information, they can lose the mental trail that lead them to a certain information item. To resolve these issues, work is being done in different visualization techniques to ease the paradigm shift.

Automated e-mail ingestion: Much of the information in CONTOUR arrived via e-mail. Future work aims to analyze the contents of these e-mails and automatically link them to the appropriate nodes in IO. Additionally, if documents are attached to the emails, they could be automatically extracted and linked into IO. This would also allow non-investigators to submit information directly in response to RFIs without requiring human intervention, simplifying the work for the investigators.

Search: Research continues into this area to determine more sophisticated means of searching through the

semantic knowledge network. One promising area is semantic search, which allows a user to search for information using the relations between nodes or pieces of information.

Rules: IO contains an inference engine that can link or create pieces of information based on a set of rules. The set of rules that will automatically create and link pieces of information in investigations is being augmented, particularly in the area of workflow.

Data Mining: As investigators accumulate data from several investigations, it may prove worthwhile to provide a means of discovering trends or similar patterns of information [5] across the investigations.

REFERENCES:

- [1] P. R. Cohen and R. Kjeldsen, "Information retrieval by constrained spreading activation in semantic networks," *Information Processing & Management*, vol23, pp. 255-68, 1987.
- [2] R. M. Keller et al., "Semantic Organizer: A Customizable Semantic Repository for Distributed NASA Project Teams," *International Semantic Web Conference (ISWC-2004)*, Hiroshima, Japan, 2004
- [3] The Comet Nucleus Tour Mishap Investigation Board, "The CONTOUR Mishap Investigation Board Report," National Aeronautics and Space Administration, Washington, DC 2003 Available: http://www.nasa.gov/pdf/52352main_contour.pdf
- [4] H. W. Gehmann et al, "Columbia Accident Investigation Board Report," Washington DC 2003 Available: http://www.nasa.gov/columbia/home/CAIB_Vol1.html
- [5] S. R. Wolfe and R. M. Keller, "Exploiting Recurring Structure in a Semantic Network," *Workshop on Mining for and from the Semantic Web, KDD 2004*, Seattle, WA, USA, 2004.